

Deep Learning with Python (taught by Dr. Christian Kauth)

Deep learning with neural networks is a fascinating field. The mixture of faster hardware, new techniques, highly optimized open source libraries and large datasets allow very large networks to be created with frightening ease. Deep neural networks have repeatedly proven impressively skillful on a range of problems.

This course is a guide to deep learning in *Python*. You will discover the *Keras* Python library for deep learning and how to use it to develop and evaluate deep learning models. You will discover the techniques and develop the skills in deep learning that you can then bring to your own machine learning projects.

After familiarizing with *Keras*, we will illustrate the skill of deep learning on some well-understood case study machine learning problems from the UCI Machine learning repository (<http://archive.ics.uci.edu/ml/index.php>) and compare the performance to the classical machine learning approaches used in the course “Introduction to Python for Predictive Modeling”. Next we introduce convolutional layers to the networks and use them to classify handwritten digits (MNIST dataset <http://yann.lecun.com/exdb/mnist/>) and real-world objects (CIFAR-10 <https://www.cs.toronto.edu/~kriz/cifar.html>). Finally, we will use deep generative models to encode images into very low dimensionality space and act on that space to tune targeted features of the image (surprise dataset).

Objectives

- To understand the structure and working principles of neural networks.
- To gain insights into some layer types of feed-forward neural networks (dense, convolutional, dropout) and how they are trained.
- To learn how to classify images with neural networks.
- To learn how to generate images from neural networks.
- To gain hands-on experience with Python and the deep learning library *Keras*.

Content

- Introduction to neural networks (structure, training methods, data preparation)
- Introduction to *Keras* (basics, model saving, checkpointing)
- Multilayer perceptrons and their performance vs. classical machine learning algorithms.
- Convolutional neural networks and their performance on image datasets (MNIST, CIFAR-10)
- Deep generative models (dimensionality reduction, principal components, autoencoders and variational auto-encoders)

Preconditions

- Basic fluency in the programming language “Python”, as e.g. provided in the course “Introduction to Python for Predictive Modeling”.
- Either a local development environment with Python, Jupyter, *Keras*, TensorFlow and admin rights, or a Google account to access Google Colab.

Duration

- 1 day on Feb 10th (roughly 7*45 minutes each day)

Evaluation

- take home exam: project work to be solved in Python

ECTS

- 0.5